

The purpose of this paper is to introduce the SmartFog cloud computing capability and highlight its ability to solve various cybersecurity challenges, including spotting anomalous cyber activities.

The Technica Innovation Platform White Paper Series presents advanced topics that will drive competitive advantage for next-generation IT over the next three-to-five years.



SMARTFOG FOR CYBERSECURITY

SmartFog is a software-based, hardware-agnostic, prototype-level capability that builds upon multiple core technologies. When employed for Cybersecurity, SmartFog leverages the convergence of the Internet of Things (IoT), Fog Computing, and Artificial Intelligence (AI) assists human cyber analysts in making sense of the glut of cyber-data. SmartFog easily deploys AI-Analytic Microservices to the Cloud, Fog, and IoT devices. The architecture of SmartFog allows for flexible integration with existing systems, offers numerous implementation options, and creates the opportunity to enhance security and privacy from Cloud-to-Things (C2T). Additionally, SmartFog Cybersecurity Microservices radically transform next-generation information security, by allowing the easily deployment, training, re-training, and re-deployment of AI-Analytic Microservices from C2T. As an example, the SmartFog Anomaly Detection Microservice utilizes an advanced machine learning neural network to spot data that is abnormal, like an atypical amount of traffic from a never-seen-before IP address. SmartFog for Cybersecurity is not intended to replace any existing solutions, but to augment existing solutions with next-generation, AI-Algorithms.

Technica has concentrated heavily on optimizing the SmartFog Cybersecurity Microservices Catalog to work on IoT devices that feature limited compute, memory, and power. Moreover, Technica has developed a Federated Learning capability that allows the SmartFog Cybersecurity Microservices to operate in Denied-Disconnected, Intermittent, Limited (D-DIL) connection environments, like those found at the tactical edge, while at the same time taking advantage of SmartFog’s architecture to enhance overall security and privacy.

IoT is the network created by connected devices—mobile phones, sensors, cameras, actuators, microcontrollers and other devices with embedded software, etc. The growth in the deployment of these devices is expected to be exponential. Already, there are more IoT devices than human beings on the planet, creating heretofore unprecedented amounts of data. IoT devices will place incredible stresses on even the most advanced Big Data infrastructures and radically increase the number of threat vectors affecting enterprise cybersecurity posture. Gartner predicts that by 2022, more than 50% of enterprise-generated data will be created and processed outside of the data center, i.e. the core or cloud¹.

¹ Gartner, “How Edge Computing Redefines Infrastructure”, Bittman, Gill and Markannen, August 2018

The Fog Computing architecture incorporated into SmartFog will assist in processing the data from IoT devices, allowing compute, network, storage, acceleration, analysis, and management functions needed for advanced cybersecurity analytics to be delivered where they are needed along a continuum of cloud to the network edge or C2T. Specifically, this means that AI algorithms like the Anomaly Detection Microservice can be deployed wherever it is needed to support a given use case—in the cloud, fog, or edge. The algorithm can be trained to spot anomalies specific to its own environment. This means that a single algorithm can serve multiple purposes, tuned for its specific situational parameters.

SMARTFOG CORE TECHNOLOGIES

Fog Computing

While enterprises have aggressively moved to deploy services to the cloud and cloud-enabled their own applications, the cloud is not best for every use case. While up-to-date network infrastructure and robust wireless connectivity can deliver request/response from device to cloud in less than 400 milliseconds, some applications need even less latency.

For example, the time from sensing to actuation on robots in the factory floor typically needs to work at 10 milliseconds or below. In some cases, the increased latency inherent to the cloud can cause sickness (in the case of Augmented Reality (AR) or Virtual Reality (VR)) or endanger human life—consider if semi-autonomous car had to talk to the cloud before engaging brakes.

Moreover, many environments that Technica supports, like the tactical edge, encounter limited or intermittent bandwidth. In such D-DIL cases, centralized applications—delivered from the cloud—do not meet mission requirements extending the Observe, Orient, Decide, and Act (OODA) Loop beyond what is acceptable.

The SmartFog prototype capability is composed of a number of cutting-edge technologies. Many are based on open source, Agile, and the DevOps (Development/Operations) movement. The architectural constructs of Fog Computing, Microservices Architecture (MSA), containerization, and DevOps are inherent or “baked-in” to the architecture of the SmartFog capability. It is important to understand these underlying software components and the architectural approaches in order to appreciate the benefits that SmartFog brings to cybersecurity solutions.

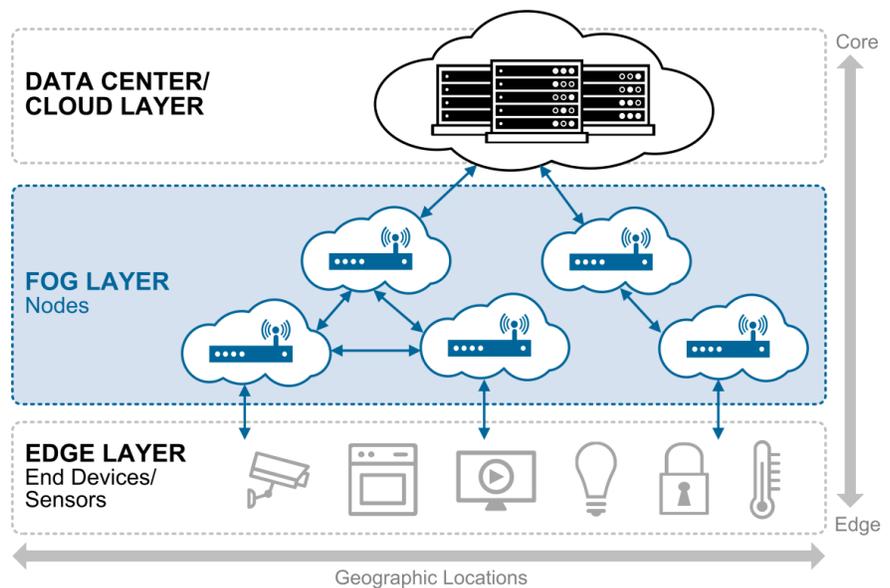


Figure 1 – Notional Fog Computing Architecture

It is important to note that the Fog Layer can be organized hierarchically. For example, consider a campus of smart buildings. Each floor in a smart building may be configured to roll up condensed data to a single, building fog node. The building fog node can be programmed to further condense data and communicate with a campus wide fog node. Finally, the campus-wide fog node can then communicate to the cloud where an enterprise could get a global view of their campus environments. Along each step up the hierarchy, microservices could perform specific functions/actions upon the data, for example sending alarms or alerts, turning on fire extinguishers, etc.

Fog Computing is an emerging computing paradigm that brings compute, storage, networking, acceleration, analytics, and management control closer to the network edge. The name comes from the fact that while the cloud is high in the air, the fog is closer to the ground. The key aspect of Fog Computing is the inclusion of a Fog Layer between the cloud (or enterprise servers in the core) and IoT devices as shown in **Figure 1**. Events from IoT devices are sent to near-by fog nodes where “intelligence” is required at or near the edge of the network.

Distributed fog nodes supplement the centralized cloud to move high volumes of data from and between edge devices, sending only the data to the cloud that is needed for longer-term analytics or management control.

The benefits of Fog Computing for IoT are numerous and can be encompassed by the acronym SCALE (Security, Cognition, Agility, Latency, and Efficiency):

- **Security**

- Like any distributed computing infrastructure, best practices like encryption of data at rest and in motion must be followed. A complete discussion of information security practices for Fog Computing is beyond the scope of this document. However, SmartFog’s Architecture offers the ability to improve overall enterprise security and privacy, especially as IoT devices are on-boarded. SmartFog for Cybersecurity is not intended to replace any existing solutions, but to enhance existing solutions and augment human cyber operators.
- Fog nodes can house sophisticated security solutions, managing all devices under their stewardship via whitelisting and encryption to ensure trusted communications with the cloud and IoT devices. This is a major benefit to IoT devices that often times lack the compute power to deal with cryptographic functions and other security enhancing applications.
- Given that not all IoT data must traverse the wire and be stored in the cloud with SmartFog’s architecture, privacy and information security can be enhanced.

- **Cognition**

- Compute can be nomadic, deployed wherever it makes sense for a given use case. This gives an enterprise greater insight and capability to respond in real-time—a continuum of intelligence. This continuum of intelligence enables the enterprise to move beyond batch-processed insights toward more real-time, event-based applications.
- Fog computing enables a next-generation AI infrastructure. AI-Analytic Microservices and be dispersed to where they are needed along the C2T continuum. This means that AI inference can occur on fog nodes or IoT devices, extending AI’s reach beyond a cloud-only capability. Likewise, AI model training can occur in the cloud on in fog nodes.

- **Agility**

- Compute, storage, networking, analytics, and control can be deployed wherever they are needed, whenever they are needed along the C2T continuum (cloud (off-premise datacenter), core (on-premise datacenter), fog nodes, and network edge). This allows for rapid innovation and affordable scaling of cybersecurity algorithms/ microservices.

- Functionality without connectivity, i.e., devices do not need an “always-on” cloud connection to perform services. For example, consider a smart-thermostat that continues to perform its functions and be programmable even when disconnected from the cloud.
- **Latency**
 - Reduced latency is one of the most oft-cited benefits of the Fog Computing architecture. Fog nodes provide low latency communications for delay sensitive processes, such as a self-driving car that cannot wait to talk to the cloud before braking is initiated or a factory floor in which robots must interact with humans in the assembly of automobiles.
 - At the tactical edge and other D-DIL environments, functionality like AI-algorithm inference and federated learning (details in a following Section) can continue to occur. The Fog Architecture can compress the ODA Loop.
- **Efficiency**
 - Architectural flexibility for specific use cases to process data where it makes the most sense. Data can be localized, i.e. operated upon where it most makes sense. It does not always need to traverse the network and reach the cloud. Additionally, fog nodes can be arranged hierarchically further extending flexibility and efficiency.
 - Lowered bandwidth requirements allow for backhaul optimization, making networks more efficient and cost effective.
 - IoT devices can consume less power and need less processing capabilities because they can rely upon compute in the fog nodes. In other words, IoT devices need only focus on their precise task like image recognition, while off-loading ancillary tasks like alerting a management console to fog nodes.

As a prototype Fog Computing platform, SmartFog incorporates all of these SCALE benefits. Specifically for cybersecurity, SmartFog enables AI algorithms packaged as Cybersecurity Microservices to perform analytic functions tailored to their environment. For example, the Anomaly Detection Microservice housed in the cloud will be trained to spot different cyber anomalies than the Anomaly Detection Microservice operating on an edge-located microcontroller facing D-DIL constraints.

Fog Computing is not a Technica derived concept. Standards bodies like the OpenFog Consortium² and the National Institute of Standards (NIST)³ have promulgated the vision of Fog Computing that is based upon open standards, fosters interoperability, and avoids vendor lock-in. Technica has closely followed these standards. Recently, IEEE adopted the IEEE 1934 standard for a Fog Computing reference architecture⁴. SmartFog conforms to this standard.

GPU Acceleration

AI usage and performance has increased dramatically in the past five years, and Deep Learning algorithms have produced the lion’s share of AI advancements. Deep Learning is a specific form of the more generalized category of machine learning. Deep Learning algorithms utilize Artificial Neural Networks (ANNs) that follow various methodologies in their constructions, e.g. Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Autoencoders, Bitwise, etc.

² <https://www.openfogconsortium.org/>

³ <https://www.nist.gov/news-events/news/2018/03/nist-releases-special-publication-500-325-fog-computing-conceptual-model>

⁴ <https://standards.ieee.org/news/2018/ieee1934-standard-fog-computing.html>

ANNs are trained to process streams of information and recognize patterns, for instance facial recognition, image recognition, speech recognition, and anomaly detection.

GPUs have drastically cut the time it takes to train the neural network model. After the model is trained, it can be deployed onto a device to perform recognition of the desired pattern, in a process called inference. Technica has created a number of AI-Analytic Microservices containing Deep Learning algorithms that can be trained to perform cybersecurity tasks.

As previously mentioned, applications that incorporate AI and IoT immensely benefit from Fog Computing. AI is about programmatically manipulating big data to recognize patterns and make real-time and time-sensitive decisions. IoT includes all the devices that are collecting, distributing, and processing data at the Edge. The compute power within the Fog Nodes allows data processing required for AI to be performed closer to the edge, without encountering the latency of the cloud. Technica's SmartFog emphasizes AI and the management of AI and system-level microservices, including versioning and configuration.

Microservices Architecture (MSA)

MSA is a specific type of software development that concentrates on building single-purpose modules with well-defined interfaces and operations. The MSA paradigm has grown in popularity in recent years as the enterprise seeks to become more agile and move towards the continuous integration/testing pattern found in DevOps solutions. Many open source projects have facilitated MSA adoption. Microservices can help create scalable, testable software that can be delivered daily/weekly.

MSA can best be thought of as the next generation of Service Oriented Architecture (SOA). The goal of SOA which exploded on the scene in the mid-2000s was to break down monolithic legacy applications into constituent services. The services would communicate, typically via SOAP/HTTP. The intent was to create lightweight, loosely coupled services; however, in practice, SOA was brittle. Any change to a service would typically break consuming applications, necessitating code changes.

MSA has the same goals as SOA, namely breaking down individual applications into lightweight services that can be re-used. However, whereas SOA essentially used middleware to glue components together, MSA uses defined Application Programming Interfaces (APIs). These APIs are typically REST based.

The SmartFog prototype is composed of a plethora of microservices performing system level functions (message brokering, database functions, complex event processing, data transformation, etc.) and AI analytic functions/algorithms (anomaly detection, facial recognition, object detection, etc.)

Containerization/Docker Containers

A container is a stand-alone unit of software that performs a specific function, also known as a microservice. Docker is an open source implementation of Linux containers. Docker operates similarly to virtualization technologies like VMWare,

Given the D-DIL environments SmartFog is designed to operate with, Technica created a lightweight message brokering microservice that utilizes MQTT (Message Queuing Telemetry Transport), although REST interfaces can be leveraged for non-D-DIL situations. Through the incorporation of MSA, SmartFog allows for the creation of scalable, stand-alone testable microservices/applications that can be delivered daily/weekly vs. yearly. Containerization is a great aid to the scalability of MSA.

but is much more lightweight in that it contains a stripped down version of Linux. A Docker container image is a stand-alone, executable software component that includes everything needed to run the microservice.

The beauty of microservices is that once constructed, and saved as a Docker image, the microservice will operate exactly the same on any piece of hardware—including IoT devices—that can run Linux. Various microservices can interoperate with one another on a single piece of hardware without worrying about versioning issues and other software incompatibility problems.

All of Technica's AI-analytic microservices are implemented as Docker containers and have well-defined APIs. They are scalable and can be architected for high-availability and fault-tolerance. The microservices can be applied to numerous use cases, including cybersecurity, and implemented at various points across an IT architecture along the C2T continuum. In addition, they may be used individually or in conjunction with an existing platform, solution or product. In other words, the Microservices Catalog developed by Technica can work apart from SmartFog.

DevOps

Microservices and containerization are key enabling elements for an Agile DevOps strategy. DevOps is a concept that severs the long bridge that typically exists between the enterprise's software development teams and operations teams in charge of software performance in a production environment. Shorter gaps between the teams equate to more updates and improvements to software functionality, and less downtime for the enterprise. A DevOps mindset is central to SmartFog because software functionality can be delivered precisely with fewer bugs, including specially trained SmartFog Cybersecurity Microservices.

SMARTFOG PROTOTYPE PLATFORM AND MICROSERVICES

Technica's Independent Research and Development (IR&D) division developed the SmartFog prototype platform to bring functions (compute, storage, networking, acceleration, analytics, and management control) closer to the edge—where the IoT devices reside.

This allows IoT events to be processed in near real-time. Importantly, SmartFog allows for data localization, i.e. data can be processed near the edge. This offloads some of the analytics burden from the cloud or core on-premise datacenters. Faster results are obtained, and with less security risk, than transmitting all data to central servers for processing.

An intuitive web-based user interface provides a full suite of services for administrators to configure a microservice, deploy it to a fog node, monitor its status, start and stop microservices, and send an updated configuration to a running microservice. These functions are initiated on individual nodes, or on clusters of nodes.

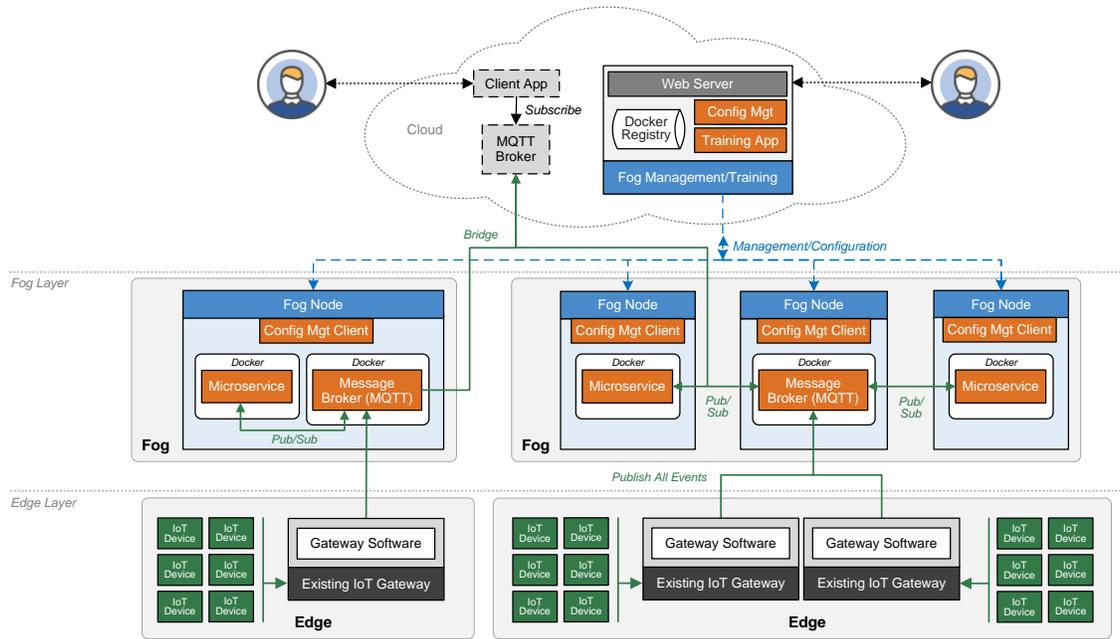


Figure 2 – Notional SmartFog Architecture

AI Analytic Microservices

- Deep Learning Algorithms
 - Fall Detection – Convolutional Neural Network (CNN)
 - Image Classification – CNN
 - Time Series Analysis – Recurrent Neural Network (RNN)
 - Long Short-Term Memory (LSTM)
 - Anomaly Detection – Autoencoder
 - Generative Adversarial Neural Network (GAN) – Technica recently began development on this microservice
- Genetic algorithms to optimize hyper-parameters
- Federated Learning

System-Level Microservices

- Complex Event Processing (CEP)
- MQTT Message Brokers (Mosquito & EMQTT)
- Machine Learning Database, e.g., perform Time Series Analysis
- Data Transformer, e.g., binary data transformed to JSON

SmartFog Architecture

Figure 2 depicts a notional SmartFog architecture. While numerous low-powered devices like Raspberry Pis can serve as fog nodes, Technica has most often utilized NVIDIA Jetson TX2s to manage the Fog Layer. These GPU accelerated devices enable IoT devices to operate with minimal compute, power, and storage capabilities, as these functions can be offloaded to the fog nodes.

Using the MQTT publish-and-subscribe protocol, SmartFog can continue to work in scenarios where network bandwidth is limited or intermittent. MQTT messages, built on top of TCP/IP, allow Fog Layer microservices to communicate with IoT devices and the cloud. Other message brokering protocols, like Advanced Message Queuing Protocol (AMQP), can be supported.

There are two types of end-users shown in the diagram. The end-user on the left represents those that employ a client application that make use of the data coming from the edge and fog. The application listens for the data, processes it in some way, and formats the data for users to see. The end-user on the right side is a SmartFog administrator, who deploys and manages the microservices from the cloud.

SmartFog Microservices Catalog

SmartFog Microservices provide discrete functions and can be viewed as analogous to apps on smartphones. However, unlike most smartphone apps, SmartFog Microservices can pass messages between themselves to create composite microservices. For example, the Anomaly Detection Microservice can communicate with the Complex Event Processor (CEP) Microservice to trigger alarms or alerts. Microservices can be dynamically updated, similar to upgrading apps on smartphones.

While SmartFog is a powerful platform, the platform is only as powerful as the jobs it can perform. Thus, the innovations involved in creating the platform, e.g., enabling Docker images on the TX2 that can be GPU accelerated, are matched with innovative microservice offerings. The current Microservice Catalog is included in the in the left-hand column.

While a microservice can be created for nearly everything— just like smartphone apps—Technica sees the greatest value in providing microservices that take advantage of hardware acceleration and offer AI capabilities with Deep Learning algorithms. The coming tsunami of IoT devices and even larger volumes of data from these devices will necessitate automated solutions that take advantage of the latest breakthroughs in AI. There will simply be too much data for humans to process. Utilizing AI in the fog layers is one of the most important benefits of Fog Computing in general and SmartFog specifically.

This is especially true in applying these advanced AI algorithms to cybersecurity. A detailed description of all entire SmartFog Microservices Catalog is beyond the scope of this document, however the following section will concentrate on the microservices most readily applicable to cybersecurity.

SMARTFOG CYBERSECURITY MICROSERVICES DESCRIPTION

The SmartFog Cybersecurity Microservices Catalog is currently composed of the following Microservices:

- **The Anomaly Detection Microservice**
- **LSTM Microservice**
- **Federated Learning Microservice**
- **Generative Adversarial Neural Network (GAN) Microservice**

Anomaly Detection Microservice (ADM) and LSTM

The Anomaly Detection Microservice (ADM) uses a specific neural network architecture—an autoencoder—to compress and decompress data as shown in **Figure 3**. The orange neural layers in the figure reduce the input into a compressed feature vector. The green neural layers attempt to expand the feature vector and recreate it. Since the network is trained to reproduce common data better, data with more decompression errors is identified as anomalous.

For cybersecurity time series analysis, Technica utilizes a Long-Short Term Memory (LSTM) neural network. LSTM is state-of-the-art for time series neural networks. Since there are several parameters that must be provided by the user to achieve optimal performance, Technica IR&D developed a genetic algorithm to select these parameters automatically.

Federated Learning Microservice

One of the core tenets of SmartFog is to enable AI to the edge. AI algorithms, like the Deep Learning based Autoencoder in the ADM and the LSTM algorithm, are typically not designed for D-DIL scenarios. Most AI algorithms are designed for a world in which connectivity is taken for granted—and memory, storage, compute, and power consumption are unconstrained. This means that most AI algorithms

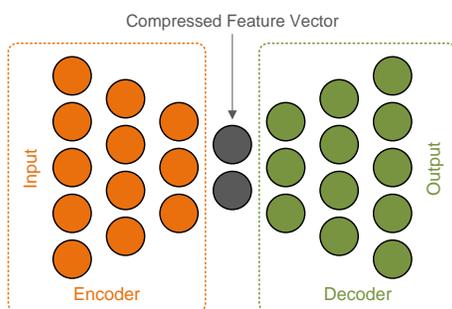


Figure 3 – ADM Autoencoder

(and therefore the latest advanced in Deep Learning) are consigned to on-premise data centers or in the cloud.

For this reason, Technica has spent much time developing a Federated Learning Microservice. This Microservice allows for algorithms, e.g. the ADM, to be deployed on multiple devices in the edge or fog nodes. Each specific instance of the ADM can be trained to spot anomalies within their current environment, even in a disconnected scenario. When connectivity is regained, the local anomaly models can be merged, updated, and distributed. In this way, all of the models can continuously learn from one another and the models can continue to detect anomalies even when disconnected from the network.

Thus, when combined with the ADM, federated learning allows for multiple anomaly detection models, each trained on their own distribution of data, to be combined into a single model that can identify anomalies across the overall distribution.

GAN Microservice

Technica recently began development of a GAN Microservice. GANs, like the Autoencoder present in ADM, are composed of two neural networks (NNs) as shown in Figure 4. A Generator NN tries to create something fake or falsified. The second NN, known as a Discriminator NN, tries to evaluate whether the Generator NN’s creation is falsified or not.

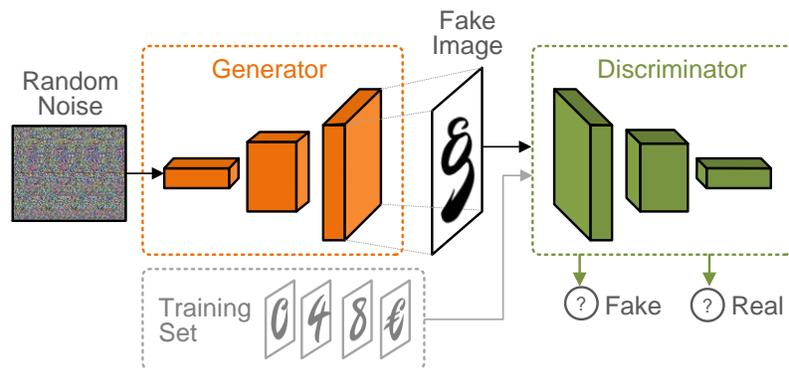


Figure 4 – GAN Architecture

Examples of use cases in which to identify anomalies are:

- Originating IP Address of Connection
- Length of Connection
- Number of Bytes Transferred
- Service Utilized (SSH, FTP, etc.)
- Computed Features, and Number of Connections to/from Same IP to Detect Specific Types of Attacks, e.g. Denial of Service (DoS)
- Determination of All Log Instances Involved in an Attack

The Discriminator NN continuously learns thereby improving its accuracy over time. Given appropriate training data, the GAN Microservice could constantly challenge a cybersecurity system's results and seek to improve defenses against increasingly sophisticated adversaries. Technica continues to perform research and development efforts to advance the GAN Microservice.

PAST PERFORMANCE

In a relatively short amount of time, SmartFog has generated interest from a diverse set of customers with varied needs. Our current qualifications include providing AI cybersecurity algorithms to the US Army Light-Weight Analytics Capability (LWAC). These cybersecurity microservices will be leveraged on the Army’s Big Data Platform (BDP) in both enterprise and tactical environments.

LWAC allows operators in the field to have access to many of the tools and capabilities provided in the current Army BDP while in a disconnected or low bandwidth environment. The ADM and LSTM Microservices with GPU accelerated AI algorithms were optimized for tactical BDP to perform advanced network analytics (anomaly detection, predictive analytics, etc.) that can be applied to cybersecurity. Network anomalies are found by close inspection of packet capture (PCAP) and log data.

Technica looks forward to incorporating the Federated Learning Microservice to bring AI (in this case anomaly detection via the autoencoder algorithm in the ADM) to the tactical edge.

Technica has also recently been selected by the Army Research Lab (ARL) Computational and Information Sciences Directorate (CISD) to further develop a prototype based on SmartFog to provide distributed processing in a heterogeneous tactical environment. CISD serves as the principal Army organization for basic and applied research in information sciences, network sciences, battlefield environment, and advanced computing and computational sciences to provide the Warfighter with knowledge superiority and ensure U.S. military superiority. During the duration of the Agreement, SmartFog will be used to advance the conceptual framework known as the “Internet of Battlefield Things” (IoBT).

CONCLUSION

Technica has been advancing cutting-edge projects at the convergence of AI, IoT—and now—Fog Computing, each directed toward transitioning our clients towards more real-time, event-driven architectures along the C2T continuum. The company believes that Fog Computing will be as disruptive a technology as cloud computing, impacting everything from computing architectures to business models. Technica has developed the SmartFog prototype capability to leverage these converging forces to improve enterprise Security and privacy, Cognition, Agility, decreasing Latency, and enhancing Efficiency (SCALE).

When applied to AI in general and cybersecurity specifically, Technica’s SmartFog enables the enterprise to take advantage of the next-generation of AI, including Deep Learning, by producing prototypes that leverage specifically tunable AI algorithms.

Currently, Technica is working with the US Army to deliver anomaly detection for network traffic and the SmartFog platform for research into the IoBT. These use cases can be expanded on an as-needed basis. Additionally, Technica continues to develop new AI algorithms, like the GANs, and package them as reusable microservices. This approach enables the cybersecurity AI algorithms to be trained to perform inference wherever they are needed along a continuum of cloud-to-fog-to edge.

Technica provides professional services, products, and innovative technology solutions to the Federal Government. We specialize in network operations and infrastructure; cyber defense and security; government application integration; systems engineering and training; and product research, deployment planning, and support.

Technica[®]

22970 Indian Creek Drive, Suite 500
Dulles, VA 20166
703.662.2000